

## Stories of Discovery &amp; Innovation

## Superfast Search Engine Speeds Past the Competition

## CONTACT US

## Office of Science

U.S. Department of Energy  
1000 Independence Ave., SW  
Washington, DC 20585  
P: (202) 586-5430  
[More Information »](#)

05.31.11

## Superfast Search Engine Speeds Past the Competition

Software developed for analyzing physics data finds powerful commercial applications.

Print SC RSS Feeds Text Size: [A](#) [A](#) [A](#)

Share / Bookmark

Our world is increasingly data-driven, whether we are searching for information on our home computer, accessing databases for everything from medical records to financial data, or scanning the depths of outer space to unlock the secrets of the universe. As data volumes increase, making sense of all this data increasingly requires an ability to quickly find essential pieces of information buried in a mountain of bytes.

To address this challenge, computer scientists at the U.S. Department of Energy's Lawrence Berkeley National Laboratory (Berkeley Lab) developed a new approach to searching massive databases. Embodied in open-source software called FastBit, the new method can search massive databases 10 to 100 times faster than large commercial database software, depending on the specific application. Originally developed to sort through the massive data produced by nuclear physics experiments, the software has found important commercial uses. Yahoo! has used FastBit to better match web content and advertisements. A German-based pharmaceutical firm has used the software to accelerate drug discovery. Still other companies have used it to analyze computer network performance or rapidly comb through masses of financial data.

FastBit was originally designed to help nuclear and high energy physicists sort through billions of data records to find a few key pieces of information. At Brookhaven National Laboratory's Relativistic Heavy Ion Collider, gold ions collide at near light speeds, generating billions of collision events, with multiple data points collected by large detectors such as the Solenoidal Tracker, or STAR. Of these billions of collisions, only a few hundred might have the distinctive signature of a new state of matter—called quark–gluon plasma—and finding evidence of this plasma has been a major objective of the STAR experiment. In such an application, the dataset includes hundreds of searchable attributes.

This kind of analysis calls for an approach fundamentally different from that of an Internet search engine or a typical commercial database.

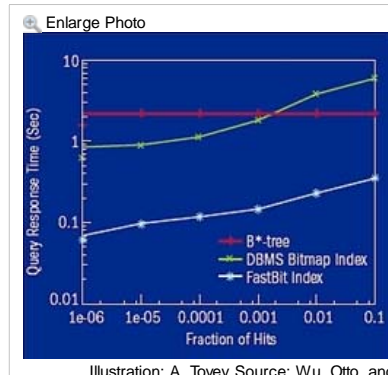

Google and Yahoo! have developed massive infrastructures to make keyword searches quick and easy. However, their search engine techniques are not appropriate for analysis tasks where most of the data records are represented by numerical values instead of text, and most search operations require full and complete answers instead of some "Top 10" records.

Commercial database systems, meanwhile, are generally designed to manage a relatively small number of searchable attributes. For example, a banking application might only be able to search for accounts based on account number and customer name. In addition, commercial database management systems are normally designed to locate an individual record (or a very small number of records) efficiently, while most scientific data analysis tasks require a much larger number of records. Furthermore, scientific data analysis requires flexibility: researchers will generally wish to examine many different scenarios or combinations of conditions and attributes.

All these requirements present formidable challenges for existing indexing methods.

The FastBit team overcame these challenges through a combination of techniques. FastBit organizes data into formats known as "Bitmap indices." Bitmap indices translate variable values into strings of bits, or 1's and 0's. Bitmap indices tend to be very efficient because computer processors are optimized to perform so-called logical operations on bits. Typically, however, bitmap indices have been used where variables have what is called low "cardinality"—that is, a limited number of possible values. Examples would include the gender or the state or zip code of the customer in the database; there are only so many genders, states, or even zip codes.

Scientific data, by contrast, typically has an enormous range of values, so further techniques for developing the index were needed. These included an alternative method for partitioning or organizing the data;



innovative ways of encoding the indices; and a revolutionary patented data compression system.

Typically, commercial database programs partition or split up data by record or groups of records. A record might include the following variables: customer name, address, phone, account balance, and date and amount of last payment. That is known as "horizontal" partitioning. (Imagine the record as a horizontal row in a spreadsheet with customer name followed by the rest of the variables.) By contrast, in the STAR application, there are billions of events (records) stored, each of which has multiple variables. But searches are usually looking for just a few variables. It would be enormously time-consuming to call up billions of whole events or records with all their variables when searching for just a few variables. So rather than partition the data by events or records, FastBit partitions data by variable—so-called "vertical" partitioning. This cuts down enormously on memory overhead and speeds processing.

In addition, FastBit provides multiple nested levels of encoding, with the top level providing a relatively coarse index to the data and each successive lower level providing finer detail. In effect, the top level indices provide pre-computed answers to anticipated queries. This enables a rapid narrowing of the search as the software zeros in from a general picture to ever more precise detail.

Finally, FastBit's authors devised an ingenious, patented method of compressing the bitmap indices that enables rapid performance of logical operations simultaneously on large swaths of data.

This combination of innovative techniques vastly improves the speed of searching operations and expands the types of data that bitmap indices can work with efficiently.

Thus a program that originated in the pure quest for scientific discovery has proved to be a powerful tool for certain commercial applications. And indeed, FastBit's creators—Kesheng "John" Wu, Arie Shoshani, and Ekow Otoo, all of Berkeley Lab—won a 2008 R & D 100 award, bestowed by R & D Magazine and often called the "Oscar of Innovation," for the commercial potential of their creation. To date, more than 5,000 users have downloaded the FastBit software, which remains freely available at the link shown below.

—Jon Bashor, Lawrence Berkeley National Laboratory, [jbashor@lbl.gov](mailto:jbashor@lbl.gov)

## Funding

Research: DOE Office of Science, Office of Advanced Scientific Computing Research, Scientific Discovery through Advanced Computing (SciDAC) Program

Facility (National Energy Research Scientific Computing Center at Lawrence Berkeley National Laboratory): DOE Office of Science, Office of Advanced Scientific Computing Research

## Publication

K. Wu, E. Otoo, and A. Shoshani, "Optimizing bitmap indices with efficient compression," *ACM Transactions on Database Systems* **31** 1 (2006)

K. Wu, E. Otoo, and A. Shoshani, "On the performance of bitmap indices for high cardinality attributes," VLDB 2004 (<http://www.vldb.org/conf/2004/RS1P2.PDF>)

## Related Link

<http://sdm.lbl.gov/fastbit>

Last modified: 5/31/2011 10:37:16 AM

[SC Jobs](#) | [Contact SC](#) | [SC Site Map](#) | [SC RSS Feeds](#) | [SC Web Policies](#)

[DOE Phone Book](#) | [DOE Employment](#) | [DOE Site Map](#)  
[DOE FOIA](#) | [DOE Privacy Policy](#) | [DOE Web Policies](#) | [DOE No Fear Act](#)  
[DOE Small Business](#) | [DOE Information Quality](#)

Shoshani 2004

**FastBit performance in comparison to alternative search systems from well-regarded commercial database management systems. Figure shows average query response time measured on the client to answer the same set of queries on a set of high-energy physics data. Bitmap indexes from both the database management system and FastBit are compressed basic bitmap indexes; the only difference is the compression method used. Across the whole range of queries, FastBit is on average 14 times faster than the commercial bitmap index. The FastBit index is 4–30 times faster than the B\*-tree.**